

AMENDMENTS TO THE SPECIFICATION

Please replace the paragraph beginning on Page 3, line 21 and ending on Page 4, line 9, with the following:

In this configuration, user workstations (or personal computers) ~~102~~ 102a, 102b, 102c, and 102n are connected to network 101 and have access to server computers 110, 120, and 130. Each user workstation is generally associated with a particular sever computer, although, in a network system that includes a storage area network, any server can provide substantially any network services for any workstation, as needed. A user, at a user workstation ~~102~~ 102a, 102b, 102c, and 102n, issues requests for operations, such as read, write, etc., which are transmitted to the associated server computer, 110, 120, or 130, which then performs the requested operation using I/O drivers 113, 123, and 133. Servers 110, 120, and 130 perform data operations on network data that is stored in disks 142 of shared storage node 140 using connections 115, 125, and 135. Each server 110, 120, and 130 has access to any network data stored at shared storage node 140, subject to policing protocol described below. The storage area network of Figure 1 includes the physical communication infrastructure and the protocols that enable server computers 110, 120, and 130 to operate with shared storage node 140.

Please replace the paragraph beginning on Page 15, line 22, through Page 16, line 4, with the following:

Figure 3 illustrates a representative configuration of a virtual storage area network of the invention. For purposes of illustration, two server computers 310 and 320 provide network services for network 301 in the example of Figure 3. However, the architecture

of the networks of the invention can be readily scaled to networks having three or more servers, examples of which are discussed below in reference to Figure 5. Network 301 also includes any number of user workstations 302 302a, 302b, 302c, and 302n, which can be personal computers or any other computing device that receives network services from servers 310 and 320.

Please replace the paragraphs beginning on Page 17, line 21, and ending on Page 18, line 8, with the following:

Figures 6 and 7 illustrate the manner in which mirror engines 317 and 327 and dedicated link 315 615 can be used to mirror data on disks 319 and 329, thereby enabling each of servers 310 and 320 to have access to all network data. As shown in Figure 6, a user of user workstation 302a causes the user workstation to issue a write operation request for writing a data block A 350 to a disk associated with network 301. As shown in Figure 6, the write operation request is transmitted through network 301 to server A 310. Since any of the servers (e.g. server A 310 or server B 320) of network 301 can process all operation requests from any user workstation, the manner in which a particular server 310 or 320 is selected to process this operation is not critical to the invention. In order to balance the load between servers 310 and 320, any desired load balancing algorithms can be implemented. Alternatively, particular servers can be assigned to particular user workstations on a default basis.

Please replace the paragraph beginning on Page 19, line 15, and ending on Page 20, line 8, with the following:

Returning now to Figure 6, I/O driver transmits the write operation request to what is perceived by the I/O driver as being a virtual shared storage node. Physically, however, the write operation request is received by mirror engine 317. The write operation request is transmitted from mirror engine 317 to disk 319, where it is executed, resulting in data A 350 being written to a particular sector or other region of the disk. In order to mirror data A to disk 329, mirror engine 317 also transmits the write operation request to a corresponding mirror engine 327 of server B 320. The write operation request is transmitted by dedicated link 315 615 or another means for communicating. Other means for communicating may include one or combination of any wireless or hard-wire communication means. If dedicated link 315 615 is used, the physical dedicated link represents one physical difference between network 301 and conventional networks. Other embodiments of the invention are capable of mirroring the data using mirror engines 317 and 327 without dedicated link 315 615. For instance, the network infrastructure of network 301 that is otherwise used to transmit data between user workstations 302 and servers 310 and 320 can be used also to transmit mirrored write operation requests from one server to another. Dedicated link 315 615, the network infrastructure of network 301 and other means for communicating represent examples of means for communicating between servers and the mass storage devices or disks thereof.

Please replace the paragraph beginning on Page 20, line 9, with the following:

In any event, mirror engine 327 receives the mirrored write operation request and transmits it to disk 329, where it is executed, resulting in data A 350 being written to disk 329. In this manner, after user workstation 302a issues the write operation request, the data associated with the write operation request is written to disk 319 and disk 329, such that both disks include mirrored copies of the same network data. It is also noted that a similar process is performed when one of the user workstations ~~302a-n~~ 302a, 302b, 302c, and 302n issues a write operation request that causes data to be deleted from a file or otherwise deleted from disk 319. In other words, if data is deleted from disk 319, the same data is deleted from disk 329, such that the same network data is mirrored and stored at both disks.

Please replace the paragraph beginning on Page 20 line 18, with the following:

Figure 7 illustrates another instance of data being mirrored between disks 319 and 329. In this case, user workstation 302n transmits a write operation request for writing data B 352 to disk 329. The write operation request is transmitted to server B 320, where it is processed by I/O driver 323 and policing protocol module 321 in a manner similar to that described above in reference to the write operation request for data A 350 of Figure 6. In Figure 7, mirror engine 327 transmits the write operation request to disk 329, where it is executed, resulting in data B being written to disk 329. In addition, mirror engine 327 transmits the write operation requests to the corresponding mirror engine 317 of server A 310 by dedicated link ~~315~~ 615 or by another means for communicating associated with network 301. Mirror engine 317 then transmits the mirrored write

operation request to disk 319, where it is executed, resulting in data B 352 being written to disk 319. Thus, Figures 6 and 7 illustrate how any server 310, 320 in network 301 is capable of causing data associated with write operation requests to be stored and mirrored at each of the disks 319, 329 of the servers in the network.

Please replace the paragraph beginning on Page 22, line 16, with the following:

After the failure of disk 319, workstations ~~302a-d~~ 302a, 302b, 302c, and 302n can continue to issue read operation requests to be processed by the virtual shared storage node through server A 310. In this example, it is assumed that workstation 302a issues a read operation request directed to data A 350. Upon the read operation request being received by server A 310, the read operation request is received by I/O driver 313 and transmitted to mirror engine 317, which, as shown at Figure 4, is within virtual shared storage node 340.

Please replace the paragraph beginning on Page 22, line 22, and ending on Page 23, line 9, with the following:

At this point, the read operation request has transmitted in the typical manner to a storage device that is perceived, from the standpoint of server A 310, as being a shared storage node. However, as mentioned above, disk 319 is not accessible and cannot service the read operation request. Accordingly, the read operation request is transmitted to server B 320 using dedicated link 615 315. The read operation request is then used to access disk 329, which has a full copy of the network data, including data A 350. Thus, network 301 is capable of seamlessly responding to inaccessibility of disk 319 by using

mirror engines 317 and 327 to redirect read operation requests that are received by server A 310. Operation of network 301 continues uninterrupted notwithstanding the failure of disk 319. Moreover, server A 310 can respond to other network operation requests, such as write operation requests, in a similar manner after the failure of disk 319 by using the virtual shared storage node.

Please replace the paragraph beginning on Page 23, line 14, with the following:

If these functional components of server A 310 become inoperable, network 301 has a secondary way of continuing to provide access to network data according to one embodiment. In this scenario, if user workstation 302a were to issue a read operation request that would otherwise be processed by server A 310, the read operation request can be serviced by server B 320, since server B 320 has access to all network data on its disk 329. For purposes of illustration, it is assumed that the read operation request issued by user workstation 302 302a is directed to data A 350. Because server A 310 is offline, server B 320 processes the read operation request. Server B 320 uses the mirrored copy of the network data stored at disk 329 to service the read operation request and thereby provide user workstation with read access to data A 350. It is noted that conventional storage area networks also enable all servers to provide read access to all network data in the case of one of the servers of the network experiencing a failure or otherwise going offline. However, unlike conventional storage area networks, the networks of the invention do not use a physical shared storage node to provide access to all network data through any server.

Please replace the paragraph beginning on Page 25, line 1, with the following:

The methods of invention illustrated in Figures 3, 4, 6 and 7 in reference to two servers can be scaled to networks having more than two servers. For instance, Figure 5 illustrates a network according to the invention having three servers, namely, server A 510 520, server B 520 and server C 530. The operation of network 501 of Figure 5 is similar to that of network 301 described above in reference to Figures 6 and 7. For instance, when server A receives a write operation request from a user workstation 502, the write operation request is processed by I/O driver 513. Policing protocol module 511 and server A 510 operate in combination with policing protocol module 521 of server B 520 and policing protocol module 531 of server C 530.